

Memoryless Control Design for Persistent Surveillance under Safety Constraints

Eduardo Arvelo, Eric Kim and Nuno C. Martins

Department of Electrical and Computer Engineering and Maryland Robotics Center

University of Maryland, College Park, MD 20742

Email: {earvelo, eskim727, nmartins}@umd.edu

Abstract—This paper deals with the design of time-invariant memoryless control policies for robots that move in a finite two-dimensional lattice and are tasked with persistent surveillance of an area in which there are forbidden regions. We model each robot as a controlled Markov chain whose state comprises its position in the lattice and the direction of motion. The goal is to find the minimum number of robots and an associated time-invariant memoryless control policy that guarantees that the largest number of states are persistently surveilled without ever visiting a forbidden state. We propose a design method that relies on a finitely parametrized convex program inspired by entropy maximization principles. Numerical examples are provided.

I. INTRODUCTION

We develop a method to design memoryless controllers for robots that move in a finite two-dimensional lattice. The goal is to achieve persistent surveillance. The term “persistent surveillance” is used to denote the task of *continuously* visiting the largest possible set of points in the lattice. In our setup, we also impose safety constraints that dictate that certain regions are *forbidden*. The forbidden regions may represent areas in which robots cannot operate (such as bodies of water) or are not allowed to visit (such as restricted airspace). The goal is to deploy the minimum number of robots equipped with a control policy that guarantees persistent surveillance of the *largest* possible set of lattice points without ever visiting a forbidden region. The memoryless strategies proposed here are applicable to miniature robots that have severe computational constraints.

The concept of persistent surveillance is similar to the concept of coverage [5], but differs from it in that the area to be surveilled must be revisited infinitely many times. Control design for persistent surveillance has been studied in [12], [13], where a semi-heuristic control policy that minimizes the time between visitations to the same region is proposed, and in [10], which proposes an algorithm for persistent surveillance of a convex polygon in the plane. These approaches, however, are not restricted to memoryless policies and do not consider safety constraints. On the implementation front, system architectures for unmanned aerial vehicles have been designed specifically for persistent surveillance purposes [2], [11].

In this paper, we model each robot as a fully-observed controlled Markov chain with finite state and control spaces. This approach, which has been successfully used in the context

of navigation and path planning (such as in [15], [17], [16], [4]), allows for the development of robust and highly scalable algorithms. Without loss of generality, we consider robots whose state is taken as its position on a finite two-dimensional lattice and direction of motion (taken from a set of four possible orientations), and limit the control space to two control actions (“forward” and “turn right”). The limitation in the control space illustrates how constrained actuation can be incorporated in our formulation. It is important to highlight, however, that the ideas described in this paper can be extended to more general dynamics and state/control spaces.

We use a recent result in [1] to compute the largest set of states that can be persistently surveilled under safety constraints, and an associated memoryless control policy. The proposed solution relies on a finitely parametrized convex program, which is highly scalable and can be efficiently solved by standard convex optimization tools, such as [8]. The approach is based on the fact that the probability mass function that maximizes the entropy under convex constraints has maximal support [6]. We also show that the minimum number of robots needed to perform persistent surveillance of the largest set of states (without ever violating the safety constraint) is equal to the number of recurrent classes of the closed loop Markov chain under the control policy computed by the proposed convex program. The recurrent classes can be found by traversing the graph of the closed loop Markov chain.

The remainder of this paper is organized as follows. Section II provides notation, basic definitions and the problem statement. The convex program that computes the maximal set of persistently surveilled states and its associated control policy is presented in Section III. Section IV provides details on computing the smallest deployment of robots necessary for maximal persistent surveillance. We discuss limiting behavior and use of additional constraints in Section V. Conclusions are provided in Section VI. Numerical examples are given throughout the paper to illustrate concepts and the proposed methodology.

II. PRELIMINARIES AND PROBLEM STATEMENTS

The following notation is used throughout the paper:

$\mathbb{X} \times \mathbb{Y}$	set of lattice positions
\mathbb{O}	set of orientations
$\mathbb{S} \stackrel{\text{def}}{=} \mathbb{X} \times \mathbb{Y} \times \mathbb{O}$	set of robot states
$\mathbb{F} \subset \mathbb{S}$	set of forbidden states
\mathbb{U}	set of control actions
S_k	state of the robot at time k
U_k	control action at time k

The state of the robot will be graphically represented as shown in Fig. 1.




Fig. 1: Graphical representation of the state of the robot. In this examples, we use $\mathbb{X} = \{1, 2, 3\}$, $\mathbb{Y} = \{1, 2\}$, and $\mathbb{O} = \{R, U, L, D\}$, where R, U, L and D represent right, up, left and down directions, respectively.

The robot's dynamics are governed by the (conditional) probability of S_{k+1} given the current state S_k and control action U_k , and are denoted as:

$$\mathcal{Q}(s^+, s, u) \stackrel{\text{def}}{=} P(S_{k+1} = s^+ \mid S_k = s, U_k = u),$$

where $s, s^+ \in \mathbb{S}, u \in \mathbb{U}$.

We denote any memoryless control policy by

$$\mathcal{K}(u, s) \stackrel{\text{def}}{=} P(U_k = u \mid S_k = s), \quad u \in \mathbb{U}, s \in \mathbb{S},$$

where $\sum_{u \in \mathbb{U}} \mathcal{K}(u, s) = 1$ for all s in \mathbb{S} . The set of all such policies is denoted as \mathbb{K} . Note that the computation of a control action may be deterministic (when $\mathcal{K}(u, s) = 1$ for a given action u) or carried out in a randomized manner, in which case the policy dictates the probabilities assigned to each control action for a given state.

Assumptions:

- Throughout the paper we assume that \mathcal{Q} is given. Hence, all quantities and sets that depend on the closed loop behavior may be indexed only by the underlying control policy \mathcal{K} .
- When multiple robots are considered, we assume that they are identical and have dynamics governed by \mathcal{Q} . In these situations, every robot executes the same control policy. Moreover, multiple robots are allowed to occupy the same position.

Given a control policy \mathcal{K} , the conditional state transition probability of the closed loop is represented as:

$$\mathcal{P}_{\mathcal{K}}(S_{k+1} = s^+ \mid S_k = s) \stackrel{\text{def}}{=} \sum_{u \in \mathbb{U}} \mathcal{Q}(s^+, s, u) \mathcal{K}(u, s).$$

We will also refer to this quantity as $\mathcal{Q}_{\mathcal{K}}(s^+, s) \stackrel{\text{def}}{=} \mathcal{P}_{\mathcal{K}}(S_{k+1} = s^+ \mid S_k = s)$.

A. Recurrence and Persistent Surveillance

A state $s \in \mathbb{S}$ is *recurrent* under a control policy \mathcal{K} if the probability of a robot revisiting state s is one, that is:

$$P_{\mathcal{K}}(S_k = s \text{ for some } k > 0 \mid S_0 = s) = 1. \quad (1)$$

We define the *set of recurrent states* $\mathbb{S}_{\mathcal{K}}^R$ under control policy \mathcal{K} as follows:

$$\mathbb{S}_{\mathcal{K}}^R \stackrel{\text{def}}{=} \left\{ s \in \mathbb{S} : (1) \text{ holds} \right\}.$$

Remark 2.1: Membership in $\mathbb{S}_{\mathcal{K}}^R$ guarantees that once a state is visited, it will be revisited infinitely many times under control policy \mathcal{K} . It does not, however, guarantee that each state in $\mathbb{S}_{\mathcal{K}}^R$ will be visited for all initial states in $\mathbb{S}_{\mathcal{K}}^R$ because $\mathbb{S}_{\mathcal{K}}^R$ may contain multiple recurrent classes. In fact, a robot will visit a certain recurrent state s with probability one if and only if it is initialized in the same recurrent class. Moreover, note that once a robot enters a recurrent class, it will never exit under control policy \mathcal{K} .

We say a state s is *persistently surveilled* under control policy \mathcal{K} and initial state $s_0 \in \mathbb{S}$ if it is recurrent under \mathcal{K} and

$$P_{\mathcal{K}}(S_k = s \text{ for some } k > 0 \mid S_0 = s_0) = 1. \quad (2)$$

If a state s is persistently surveilled under control policy \mathcal{K} and initial state $s_0 \in \mathbb{S}_{\mathcal{K}}^R$, then it must be that s and s_0 belong to the same recurrent class.

We define the *set of persistently surveilled states* $\mathbb{S}_{s_0, \mathcal{K}}^{ps}$ under control policy \mathcal{K} and initial state $s_0 \in \mathbb{S}$ to be:

$$\mathbb{S}_{s_0, \mathcal{K}}^{ps} \stackrel{\text{def}}{=} \left\{ s \in \mathbb{S}_{\mathcal{K}}^R : (2) \text{ holds} \right\}.$$

The set $\mathbb{S}_{s_0, \mathcal{K}}^{ps}$ is a recurrent class of the closed loop dynamics $\mathcal{Q}_{\mathcal{K}}$. Note that for every state s in $\mathbb{S}_{s_0, \mathcal{K}}^{ps}$, it holds that $\mathbb{S}_{s, \mathcal{K}}^{ps} = \mathbb{S}_{s_0, \mathcal{K}}^{ps}$. Moreover, if there exists a recurrent state for which $\mathbb{S}_{s_0, \mathcal{K}}^{ps} = \mathbb{S}_{\mathcal{K}}^R$, the set $\mathbb{S}_{\mathcal{K}}^R$ has *only one* recurrent class.

Given a set \mathbb{F} of forbidden states, we define the set of states that are recurrent and for which the probability of transitioning into \mathbb{F} is zero.

The *set of \mathbb{F} -safe recurrent states* $\mathbb{S}_{\mathcal{K}, \mathbb{F}}^R$ under a control policy \mathcal{K} is defined as:

$$\mathbb{S}_{\mathcal{K}, \mathbb{F}}^R \stackrel{\text{def}}{=} \left\{ s \in \mathbb{S}_{\mathcal{K}}^R : \mathcal{Q}_{\mathcal{K}}(s^+, s) = 0, s^+ \in \mathbb{F} \right\}.$$

We define the **maximal** set of \mathbb{F} -safe recurrent states as:

$$\mathbb{S}_{\mathbb{F}}^R \stackrel{\text{def}}{=} \bigcup_{\mathcal{K} \in \mathbb{K}} \mathbb{S}_{\mathcal{K}, \mathbb{F}}^R.$$

Finally, the *set of \mathbb{F} -safe persistently surveilled states* $\mathbb{S}_{s_0, \mathcal{K}, \mathbb{F}}^{ps}$ under a control policy \mathcal{K} and initial state $s_0 \in \mathbb{S}$ is defined as:

$$\mathbb{S}_{s_0, \mathcal{K}, \mathbb{F}}^{ps} \stackrel{\text{def}}{=} \left\{ s \in \mathbb{S}_{s_0, \mathcal{K}}^{ps} : \mathcal{Q}_{\mathcal{K}}(s^+, s) = 0, s^+ \in \mathbb{F} \right\}.$$

Remark 2.2: As before, $\mathbb{S}_{s_0, \mathcal{K}, \mathbb{F}}^{ps}$ is a (safe) recurrent class of $\mathcal{Q}_{\mathcal{K}}$.

B. Problem Statement

We start by addressing the following problem:

Problem 2.3: (Maximal set of \mathbb{F} -safe recurrent states).

Given a set of forbidden states \mathbb{F} , determine:

- (a) $\mathbb{S}_{\mathbb{F}}^R$; and
- (b) a control policy \mathcal{K}^* such that $\mathbb{S}_{\mathcal{K}^*}^R = \mathbb{S}_{\mathbb{F}}^R$.

In light of Remark 2.1, note that in order to persistently surveil all possible states, we need to determine how many robots to use and in which state they should be initialized. The following problem addresses this issue.

Problem 2.4: (Maximal \mathbb{F} -safe persistent surveillance).

Given a set of forbidden states \mathbb{F} , determine the *minimum* number of robots r , a control policy $\hat{\mathcal{K}}$ and a set of initial states $\{s^1, \dots, s^r\}$, r , so that

$$\bigcup_{i=1}^r \mathbb{S}_{s^i, \hat{\mathcal{K}}, \mathbb{F}}^{ps} = \mathbb{S}_{\mathbb{F}}^R. \quad (3)$$

Remark 2.5: The following is a list of important comments on Problems 2.3 and 2.4.

- There is no \mathcal{K} such that the states in $\mathbb{S} \setminus \mathbb{S}_{\mathbb{F}}^R$ can be \mathbb{F} -safe and recurrent
- Once r robots are initialized with initial states $\{s^1, \dots, s^r\}$, it is guaranteed that the largest possible set of states will be visited infinitely many times without ever visiting a forbidden state.

We will propose a convex optimization problem that efficiently computes $\mathbb{S}_{\mathbb{F}}^R$ and a control policy \mathcal{K}^* such that $\mathbb{S}_{\mathcal{K}^*}^R = \mathbb{S}_{\mathbb{F}}^R$. We will show that the minimum number of robots r required to persistently surveil $\mathbb{S}_{\mathbb{F}}^R$ is the number of distinct recurrent classes of the closed loop Markov chain under the computed control policy \mathcal{K}^* .

III. COMPUTING THE MAXIMAL SET OF RECURRENT STATES: A CONVEX APPROACH

Let $\mathbb{P}_{\mathbb{S}\mathbb{U}}$ be the set of all probability mass functions (pmfs) with support in $\mathbb{S} \times \mathbb{U}$, and consider the following convex optimization program:

$$f_{\mathbb{S}\mathbb{U}}^* = \arg \max_{f_{\mathbb{S}\mathbb{U}} \in \mathbb{P}_{\mathbb{S}\mathbb{U}}} \mathcal{H}(f_{\mathbb{S}\mathbb{U}}) \quad (4)$$

subject to:

$$\sum_{u^+ \in \mathbb{U}} f_{\mathbb{S}\mathbb{U}}(s^+, u^+) = \sum_{s \in \mathbb{S}, u \in \mathbb{U}} \mathcal{Q}(s^+, s, u) f_{\mathbb{S}\mathbb{U}}(s, u) \quad (5)$$

$$\sum_{u \in \mathbb{U}} f_{\mathbb{S}\mathbb{U}}(s, u) = 0, \quad s \in \mathbb{F} \quad (6)$$

where $\mathcal{H} : \mathbb{P}_{\mathbb{S}\mathbb{U}} \rightarrow \mathbb{R}_{\geq 0}$ is the entropy of $f_{\mathbb{S}\mathbb{U}}$, and is given by

$$\mathcal{H}(f_{\mathbb{S}\mathbb{U}}) = - \sum_{u \in \mathbb{U}} \sum_{s \in \mathbb{S}} f_{\mathbb{S}\mathbb{U}}(s, u) \ln(f_{\mathbb{S}\mathbb{U}}(s, u))$$

where we adopt the standard convention that $0 \ln(0) = 0$.

The following proposition, which has been modified from Theorem 3.1 in [1], provides a solution to Problem 2.3.

Proposition 3.1: Let \mathbb{F} be given, assume that (4)-(6) is feasible, and that $f_{\mathbb{S}\mathbb{U}}^*$ is the optimal solution. In addition, adopt

the marginal pmf $f_{\mathbb{S}}^*(s) = \sum_{u \in \mathbb{U}} f_{\mathbb{S}\mathbb{U}}^*(s, u)$ and consider that $\mathcal{G} : \mathbb{U} \times \mathbb{S} \rightarrow [0, 1]$ is any function satisfying $\sum_{u \in \mathbb{U}} \mathcal{G}(u, s) = 1$ for all s in \mathbb{S} . The following holds:

- (a) $\mathbb{S}_{\mathbb{F}}^R = \mathbb{W}_{f_{\mathbb{S}}^*}$
- (b) $\mathbb{S}_{\mathcal{K}^*}^R = \mathbb{S}_{\mathbb{F}}^R$ for \mathcal{K}^* given by:

$$\mathcal{K}^*(u, s) = \begin{cases} \frac{f_{\mathbb{S}\mathbb{U}}^*(s, u)}{f_{\mathbb{S}}^*(s)}, & s \in \mathbb{W}_{f_{\mathbb{S}}^*} \\ \mathcal{G}(u, s), & \text{otherwise} \end{cases}, \quad (u, s) \in \mathbb{U} \times \mathbb{S} \quad (7)$$

where $\mathbb{W}_{f_{\mathbb{S}}^*}$ is the support of $f_{\mathbb{S}}^*$ and is given by $\mathbb{W}_{f_{\mathbb{S}}^*} = \{s \in \mathbb{S} : f_{\mathbb{S}}^*(s) > 0\}$.

Comments on the proof of Proposition 3.1: The proof of Proposition 3.1 closely follows the proof of Theorem 3.1 in [1] and is omitted. However, it is important to highlight that constraint (5) enforces recurrence and constraint (6) enforces \mathbb{F} -safety. Moreover, note that the pmf that maximizes the entropy under convex constraints has maximal support (see Lemma 3.5 in [1]).

Example 3.2: Let $\mathbb{X} = \mathbb{Y} = \{1, \dots, 5\}$, $\mathbb{O} = \{R, U, L, D\}$ and consider a robot whose action space is given by $\mathbb{U} = \{\text{"forward"}, \text{"turn right"}\}$. Moreover, let the set of forbidden states be given by: $\mathbb{F} = \{(x, y, \theta) \in \mathbb{S} : (x, y) \in \{(1, 1), (1, 5), (5, 1), (5, 5), (3, 3)\}\}$, which means the robot is prohibited from visiting the center or corner locations of the lattice.

In order to specify \mathcal{Q} , we first define an auxiliary conditional pmf \mathcal{Q}' defined on $\mathbb{X}' = \mathbb{Y}' = \{1, 2, 3\}$ and $\mathbb{O}' = \{R, U, L, D\}$. For clarity, \mathcal{Q}' is shown graphically in Fig. 2, which contains the probabilities of transitioning to states shown as dark triangles given the previous state shown as a white triangle. There is uncertainty only for transitions that occur on the edge of the lattice. Since we consider dynamics that are spatially invariant, the transition probabilities for states not shown in Fig. 2 can be computed by appropriate manipulation of the ones shown. Similarly, \mathcal{Q} is constructed by appropriate expansion of \mathcal{Q}' .

We use [8] to solve (4)-(6) and use Proposition 3.1 to compute $\mathbb{S}_{\mathbb{F}}^R$ and a control policy \mathcal{K}^* such that $\mathbb{S}_{\mathcal{K}^*}^R = \mathbb{S}_{\mathbb{F}}^R$. The set $\mathbb{S}_{\mathbb{F}}^R$ can be seen in Fig. 3, where the areas in red represent the states in \mathbb{F} , and the triangles in blue represent the states in $\mathbb{S}_{\mathbb{F}}^R$. The control \mathcal{K}^* , computed using (7), has been omitted due to space constraints.

IV. MAXIMAL PERSISTENT SURVEILLANCE AND ROBOT DEPLOYMENT

In this section, we provide a solution to Problem 2.4, which seeks the minimum number r of robots, a control policy $\hat{\mathcal{K}}$ and a set of initial states $\{s^1, \dots, s^r\}$ so that $\bigcup_{i=1}^r \mathbb{S}_{s^i, \hat{\mathcal{K}}, \mathbb{F}}^{ps} = \mathbb{S}_{\mathbb{F}}^R$.

In light of a previous remark, recall that the set of \mathbb{F} -safe persistently surveilled states $\mathbb{S}_{s_0, \mathcal{K}, \mathbb{F}}^{ps}$ is a *recurrent class* of $\mathcal{Q}_{\mathcal{K}}$. In practice, this means that a robot with initial state $S_0 = s_0$ applies control policy \mathcal{K} , it is guaranteed that:

- the robot will never leave $\mathbb{S}_{s_0, \mathcal{K}, \mathbb{F}}^{ps}$;
- every state in $\mathbb{S}_{s_0, \mathcal{K}, \mathbb{F}}^{ps}$ will be visited infinitely many times;
- states in \mathbb{F} will never be visited.




Fig. 2: Graphical representation of some transitions in \mathcal{Q}' .




Fig. 3: Depiction of $\mathbb{S}_{\mathbb{F}}^R$ in blue. The red areas represent the forbidden states.

To find all the (safe) recurrent classes in $\mathbb{S}_{\mathcal{K}, \mathbb{F}}^R$, flood-fill-type algorithms may be used, where the graph of $\mathcal{Q}_{\mathcal{K}}$ is traversed, either in a depth-first or breath-first manner. An edge from s to s^+ of the graph of $\mathcal{Q}_{\mathcal{K}}$ exists if and only if $\mathcal{Q}_{\mathcal{K}}(s^+, s) > 0$ holds. Note that states in $\mathbb{S} \setminus (\mathbb{S}_{\mathcal{K}, \mathbb{F}}^R \cup \mathbb{F})$ do not need to be searched.

Given \mathbb{F} and a control policy \mathcal{K} , let $n_{\mathcal{K}}$ be the number of

distinct recurrent classes of $\mathcal{Q}_{\mathcal{K}}$, and note that the following holds:

$$\bigcup_{i=1}^{n_{\mathcal{K}}} \mathbb{S}_{s^i, \mathcal{K}, \mathbb{F}}^{ps} = \mathbb{S}_{\mathcal{K}, \mathbb{F}}^R,$$

where $\{s^1, \dots, s^{n_{\mathcal{K}}}\}$ is a set of initial states, and $\{\mathbb{S}_{s^i, \mathcal{K}, \mathbb{F}}^{ps}\}_{i=1}^{n_{\mathcal{K}}}$ are distinct recurrent classes.

We define the set of all admissible control policies whose \mathbb{F} -safe set of recurrent states are maximal to be:

$$\mathbb{K}_{\mathbb{F}}^R = \{\mathcal{K} \in \mathbb{K} : \mathbb{S}_{\mathcal{K}, \mathbb{F}}^R = \mathbb{S}_{\mathbb{F}}^R\},$$

and note that in order to solve Problem 2.4, we must:

- find a control policy $\hat{\mathcal{K}}$ in $\mathbb{K}_{\mathbb{F}}^R$ such that $n_{\hat{\mathcal{K}}} \leq n_{\mathcal{K}}$ for all \mathcal{K} in $\mathbb{K}_{\mathbb{F}}^R$. Note that $n_{\hat{\mathcal{K}}}$ is the minimum number of robots needed for maximal persistent surveillance.
- identify the recurrent classes in $\mathbb{S}_{\hat{\mathcal{K}}, \mathbb{F}}^R$ (by exploring the graph of $\mathcal{Q}_{\hat{\mathcal{K}}}$); and
- select one (any) state from each of the recurrent classes to compose the set of initial states

Note that the control policy \mathcal{K}^* given in (7) is a candidate for maximal persistent surveillance since $\mathbb{S}_{\mathcal{K}^*, \mathbb{F}}^R = \mathbb{S}_{\mathbb{F}}^R$. The following proposition will show that $n_{\mathcal{K}^*} \leq n_{\mathcal{K}}$ for all \mathcal{K} in $\mathbb{K}_{\mathbb{F}}^R$.

Proposition 4.1: Let \mathbb{F} be given, and take \mathcal{K}^* to be the control policy in (7). The following holds:

$$n_{\mathcal{K}^*} \leq n_{\mathcal{K}}, \quad \mathcal{K} \in \mathbb{K}_{\mathbb{F}}^R.$$

Proof: Suppose there exists a control policy $\bar{\mathcal{K}}$ in $\mathbb{K}_{\mathbb{F}}^R$ such that $n_{\bar{\mathcal{K}}} < n_{\mathcal{K}^*}$. Since $\bar{\mathcal{K}}$ belongs to $\mathbb{K}_{\mathbb{F}}^R$, there must exist a control policy $\tilde{\mathcal{K}}$ with the same sparsity pattern as $\bar{\mathcal{K}}$ and a pmf \tilde{f}_{SU} in \mathbb{P}_{SU} that satisfies (5) and (6) for which:

$$\tilde{\mathcal{K}}(u, s) = \begin{cases} \frac{\tilde{f}_{SU}(s, u)}{\tilde{f}_S(s)}, & s \in \mathbb{S}_{\mathbb{F}}^R \\ \mathcal{G}(u, s), & \text{otherwise} \end{cases}, \quad (u, s) \in \mathbb{U} \times \mathbb{S}$$

where $\tilde{f}_S(s) = \sum_{u \in \mathbb{U}} \tilde{f}_{SU}(s, u)$ and $\mathcal{G} : \mathbb{U} \times \mathbb{S} \rightarrow [0, 1]$.

Since $\mathcal{Q}_{\bar{\mathcal{K}}}$ has fewer recurrent classes than $\mathcal{Q}_{\mathcal{K}^*}$, there must exist a pair (s, u) in $\mathbb{S}_{\mathbb{F}}^R \times \mathbb{U}$ for which $\bar{\mathcal{K}}(s, u) > 0$ and $\mathcal{K}^*(s, u) = 0$ holds. Since $\bar{\mathcal{K}}$ and $\tilde{\mathcal{K}}$ have the same sparsity pattern, it holds that $\tilde{\mathcal{K}}(s, u) > 0$. Therefore, it must be that $\tilde{f}_{SU}(s, u) > 0$ and $\tilde{f}_S(s) > 0$. In other words, the support of \tilde{f}_{SU} is not contained in the support of \tilde{f}_S , which is a contradiction by Lemma 3.5 in [1]. ■

Remark 4.2: Suppose we change the objective function in (4) to $\mathcal{H}(f_S)$ and add the following constraint: $f_S(s) = \sum_{u \in \mathbb{U}} f_{SU}(u, s)$. Note that an appropriate modification of Proposition 3.1 would enable us to find $\mathbb{S}_{\mathbb{F}}^R$ and an associated control policy (i.e., solve Problem 3.1), with the added benefit that the modified convex program would be computationally less intensive (since fewer calls to the entropy function would be required). However, maximizing the entropy of the marginal distribution f_S would not solve the problem of maximal persistent surveillance since Proposition 4.1 would not apply.

Example 4.3: Consider again the example described in Example 3.2. By exploring the graph of $\mathcal{Q}_{\mathcal{K}^*}$, we conclude

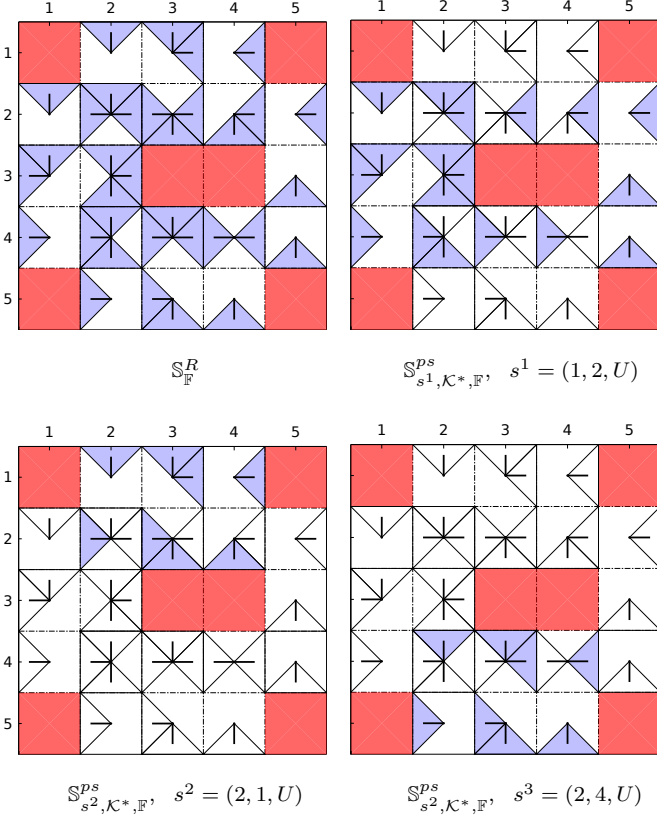


Fig. 4: Top left: maximal set of recurrent states $\mathbb{S}_{\mathbb{F}}^R$ (in blue). Others: three recurrent classes whose union is $\mathbb{S}_{\mathbb{F}}^R$.

that only one robot is required to perform maximal persistent surveillance (i.e., $\mathbb{S}_{\mathbb{F}}^R$ contains only one recurrent class). Any state in $\mathbb{S}_{\mathbb{F}}^R$ may be selected as the robot's initial state.

Suppose that we now change the set of forbidden states to include location (4, 3) (i.e., let $\mathbb{F} = \{(x, y, \theta) \in \mathbb{S} : (x, y) \in \{(1, 1), (1, 5), (5, 1), (5, 5), (3, 3), (4, 3)\}\}$). Resolving (4)-(6), applying Propositions 3.1 and 4.1, and searching the graph of the closed loop Markov chain, we conclude that at least three robots are required to perform maximal persistent surveillance of $\mathbb{S}_{\mathbb{F}}^R$ (see Fig. 4). Any state from each recurrent class may be used as initial states, so we can choose the set of initial states to be: $\{(1, 2, U), (2, 1, U), (2, 4, U)\}$. Note that the set $\mathbb{S}_{\mathbb{F}}^R$ is now smaller than in the previous example (34 vs. 40 states).

V. LIMITING BEHAVIOR AND OTHER CONSTRAINTS

We define $\mathcal{T}_{\mathcal{K}}$, the long term proportion of time the robot, under control policy \mathcal{K} , visits state s in \mathbb{S} having started at state s_0 , to be:

$$\mathcal{T}_{\mathcal{K}}(s, s_0) \stackrel{\text{def}}{=} \lim_{k \rightarrow \infty} \frac{1}{k} \sum_{i=1}^k \mathcal{I}(S_i = s, S_0 = s_0),$$

where \mathcal{I} is the indicator function.

A. Limiting Behavior with One Recurrent Class

Given a forbidden set \mathbb{F} , and let f_{SU}^* be the optimal solution to (4)-(6), and \mathcal{K}^* be the control policy computed in (7), and suppose $\mathbb{S}_{\mathcal{K}^*, \mathbb{F}}^R$ has only one recurrent class. For any initial state s_0 in $\mathbb{S}_{\mathbb{F}}^R$, the following holds with probability one:

$$\mathcal{T}_{\mathcal{K}^*}(s, s_0) = f_{\mathbb{S}}^*(s), \quad (8)$$

where $f_{\mathbb{S}}^*(s) = \sum_{u \in \mathbb{U}} f_{SU}^*(s, u)$. Since we have not imposed aperiodicity on $\mathcal{Q}_{\mathcal{K}^*}$, we cannot state stronger convergence. However, equation (8) still tells us valuable information regarding the limiting behavior of the robot.

Note that the pmf that maximizes the entropy is “as uniform as possible” (in fact, when unconstrained, the pmf that maximizes the entropy is uniform.). However, additional convex constraints can be added to our formulation in order to shape the distribution of the optimal pmf and, thus, influence the limiting behavior of the robot.

Consider the following constraint:

$$\sum_{(x,y) \in \mathbb{D}, \theta \in \mathbb{O}, u \in \mathbb{U}} f_{SU}((x, y, \theta), u) > \alpha, \quad (9)$$

where $\mathbb{D} \subset \mathbb{X} \times \mathbb{Y}$ is a region of the lattice. The set \mathbb{D} can be interpreted as a region of high interest that should be surveilled more often. Suppose the convex program (4)-(6) and (9) is feasible, that f_{SU}^{**} is the optimal solution and \mathcal{K}^{**} is the associated control policy. The following holds for any s_0 in $\mathbb{S}_{\mathbb{F}}^R$ with probability one:

$$\sum_{(x,y) \in \mathbb{D}, \theta \in \mathbb{O}} \mathcal{T}_{\mathcal{K}^{**}}((x, y, \theta), s_0) > \alpha.$$

Example 5.1: Let $\mathbb{X} = \mathbb{Y} = \{1, \dots, 10\}$, $\mathbb{O} = \{R, U, L, D\}$, and consider again a robot whose action space is given by $\mathbb{U} = \{\text{“Forward”}, \text{“Turn Right”}\}$. The dynamics \mathcal{Q} are similar to what was used in Examples 3.2 and 4.3, except that we add uncertainty to the transition of states that lie in the interior of the grid (see Fig. 5). The probabilities for states on the edge of the grid are the same as before (see Fig. 2).




Fig. 5: Graphical representation of some transitions in \mathcal{Q}' .

Consider the set of forbidden state $\mathbb{F} = \{(x, y, \theta) \in \mathbb{S} : (x, y) \in \{(2, 2), (2, 3), (3, 2), (3, 3), (8, 8), (8, 9), (9, 8), (9, 9)\}\}$. We solve (4)-(6) using the tool in [8]. In Fig. 6, each state that belongs in $\mathbb{S}_{\mathbb{F}}^R$ is shown in blue, where the darker the




Fig. 6: Depiction of $\mathbb{S}_{\mathbb{F}}^R$ in blue. Darker blue indicates a higher value for f_S^* .




Fig. 7: Depiction of $\mathbb{S}_{\mathbb{F}}^R$ in blue with additional constraint (9). Darker blue indicates a higher value for f_S^* .

blue, the higher the value of f_S^* . Note that the distribution is relatively uniform.

Consider now $\mathbb{D} = \{(x, y) \in \mathbb{X} \times \mathbb{Y} : 3 \leq x, y \leq 8\}$, and let $\alpha = 0.75$. We solve (4)-(6) and (9). The result can be seen in Fig. 7.

B. Limiting Behavior with Multiple Recurrent Classes

Consider again f_{SU}^* and \mathcal{K}^* as before, and, without loss of generality, let $\mathbb{S}_{\mathcal{K}^*, \mathbb{F}}^R$ have two recurrent classes with initial states s^1 and s^2 (i.e., $\mathbb{S}_{s^1, \mathcal{K}^*, \mathbb{F}}^{ps} \cup \mathbb{S}_{s^2, \mathcal{K}^*, \mathbb{F}}^{ps} = \mathbb{S}_{\mathcal{K}^*, \mathbb{F}}^R$). For any initial state s_0 in $\mathbb{S}_{s^1, \mathcal{K}^*, \mathbb{F}}^{ps}$ (equiv., $\mathbb{S}_{s^2, \mathcal{K}^*, \mathbb{F}}^{ps}$), the following holds with probability one:

$$\mathcal{T}_{\mathcal{K}^*}(s, s_0) = \frac{f_S^*(s)}{\beta}, \quad (10)$$

where $\beta = \sum_{s \in \mathbb{S}_{s^1, \mathcal{K}^*, \mathbb{F}}^{ps}} f_S^*(s)$ (equiv. $\beta = \sum_{s \in \mathbb{S}_{s^2, \mathcal{K}^*, \mathbb{F}}^{ps}} f_S^*(s)$).

With equation (10) in mind, note that additional convex constraints may also be used to influence the limiting behavior of the robots. Moreover, by carefully selecting the number of robots allocated to each recurrent class, one can achieve a desirable limiting behavior for the ensemble of robots.

VI. CONCLUSIONS

We have proposed methods to design memoryless strategies for controlled Markov chains that guarantee maximal persistent surveillance properties under safety constraints. The uncomplicated structure of the resulting controllers makes them implementable in small robots. We have described a finitely parametrized convex program that solves this problem via entropy maximization principles, and we show that the computed control policy results in the closed loop Markov chain with the least number of recurrent classes.

ACKNOWLEDGMENT

This work was supported by NSF Grant CNS 0931878, AFOSR Grant FA95501110182, ONR UMD-AppEl Center and the Multiscale Systems Center, one of six research centers funded under the Focus Center Research Program.

REFERENCES

- [1] E. Arvelo and N. C. Martins. Control design for Markov chains under safety constraints: a convex approach, 2012. arXiv:1209.2883.
- [2] B. Bethke, J. Redding, J. P. How, M. A. Vavrana, and J. Vian. Agent capability in persistent mission planning using approximate dynamic programming. *American Control Conference*, pages 1623–1628, 2010.
- [3] V.S. Borkar. Controlled Markov chains with constraints. *Sadhana*, 15(4-5):405–413, December 1990.
- [4] J. Burlet, O. Aycard, and T. Fraichard. Robust motion planning using Markov decision processes and quadtree decomposition. In *IEEE International Conference on Robotics and Automation*, pages 2820–2825, 2004.
- [5] H. Choset. Coverage for robotics – A survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31:113–126, 2001.
- [6] I. Csiszar and P.C. Shields. Information Theory and Statistics: A Tutorial. *Foundations and Trends in Communications and Information Theory*, 1(4):1–116, 2004.
- [7] B. Fox. Markov Renewal Programming by Linear Fractional Programming. *SIAM Journal on Applied Mathematics*, 14(6):1418–1432, November 1966.
- [8] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx/>, April 2011.
- [9] O. Hernandez-Lerma and J.B. Lasserre. *Discrete-time Markov control processes: Basic optimality criteria*, volume 30 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin and New York, 1996.
- [10] P. F. Hokayem, D. Stipanovic, and M. W. Spong. On persistent coverage control. In *IEEE Conference on Decision and Control*, pages 6130–6135, 2007.
- [11] N. Michael, E. Stump, and K. Mohta. Persistent surveillance with a team of MAVs. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2708–2714, 2011.
- [12] N. Nigam, S. Bieniawski, I. Kroo, and J. Vian. Control of Multiple UAVs for Persistent Surveillance: Algorithm and Flight Test Results. *IEEE Transactions on Control Systems Technology*, 20(5):1236–1251, 2012.
- [13] N. Nigam and I. Kroo. Persistent Surveillance Using Multiple Unmanned Air Vehicles. In *IEEE Aerospace Conference*, pages 1–14, 2008.
- [14] M.L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, New York, NY, 1994.
- [15] R. Simmons and S. Koenig. Probabilistic robot navigation in partially observable environments. *International Joint Conference on Artificial Intelligence*, 1995.

- [16] G. Theodorou and S. Mahadevan. Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292). In *IEEE International Conference on Robotics and Automation*, pages 1347–1352. IEEE, 2002.
- [17] A. Undurti and J. P. How. A decentralized approach to multi-agent planning in the presence of constraints and uncertainty. In *IEEE International Conference on Robotics and Automation*, pages 2534–2539, 2011.
- [18] P. Wolfe and G.B. Dantzig. Linear Programming in a Markov Chain. *Operations Research*, 10(5):702–710, October 1962.